Price Manipulation in the Bitcoin Ecosystem

Neil Gandal
Berglas School of Economics
Tel Aviv University, Israel
gandal@post.tau.ac.il

JT Hamrick
Tandy School of Computer Science
The University of Tulsa, USA
jth563@utulsa.edu

Tyler Moore
Tandy School of Computer Science
The University of Tulsa, USA
tyler-moore@utulsa.edu

Tali Oberman
Berglas School of Economics
Tel Aviv University, Israel
otalika@yahoo.com

Abstract

We identify and analyze the impact of suspicious trading activity (STA) on the Mt. Gox Bitcoin currency exchange between February and November 2013. We discuss two distinct STA periods in which approximately 600,000 bitcoins (BTC) valued at $188 million were acquired by agents who did not pay for the bitcoins. During the second period, the USD-BTC exchange rate rose by an average of $20 at Mt. Gox on days when suspicious trades took place, compared to a slight decline on days without suspicious activity. Based on rigorous analysis with extensive robustness checks, we conclude that the suspicious trading activity caused the unprecedented spike in the USD-BTC exchange rate in late 2013, when the rate jumped from around $150 to more than $1,000 in two months.

1 Introduction

Bitcoin has experienced a meteoric rise in popularity since its introduction in 2009 [17]. While digital currencies were proposed as early as the 1980s, bitcoin was the first to catch on. The total value of all bitcoins in circulation today is around $28 billion [6], and it has inspired scores of competing cryptocurrencies that follow a similar design. Bitcoin and most other cryptocurrencies do not require a central authority to validate and settle transactions. Instead, these currencies use only cryptography (and an internal incentive system) to control transactions, manage the supply, and prevent fraud. Payments are validated by a decentral-

Workshop on the Economics of Information Security (WEIS) 2017.
The authors gratefully acknowledge support from a research grant from the Interdisciplinary Cyber Research Center at Tel Aviv University. We thank Nittai Bergman for helpful suggestions and comments.

ized network. Once confirmed, all transactions are stored digitally and recorded in a public "blockchain," which can be thought of as an accounting system[1].

While bitcoin shows great promise to disrupt existing payment systems through innovations in its technical design, the Bitcoin ecosystem[2] has been a frequent target of attacks by financially-motivated criminals. In this paper, we leverage a unique and very detailed data set to examine suspicious trading activity that occurred over a ten-month period in 2013 on Mt. Gox, the leading Bitcoin currency exchange at the time. We first quantify the extent of the suspicious/fraudulent trading activity and show that it constitutes a large fraction of trading on the days the activity occurred. We then examine whether and how this trading activity impacted Mt. Gox and the broader Bitcoin ecosystem.

Our main results are as follows:

- Prices rose on approximately 80 percent of the days that the suspicious trading activity occurred. By contrast, prices rose on approximately 55 percent of the days in which no suspicious trading activity occurred.

- During days the key actor was active, on average, the USD/BTC exchange rate increased by more than $20 a day. During the same period, the exchange rate was virtually unchanged on the days in which the actor was not active.

- The suspicious trading activity of a single actor was the primary cause of the massive spike in the USD/BTC exchange rate in which the rate rose from around $150 to over $1,000 in just two months in late 2013. The fall was nearly as precipitous: the Mt. Gox exchange folded due to insolvency in early 2014 and it has taken more than three years for bitcoin to match the rise triggered by fraudulent transactions.

1.1 Why Should We Care?

Why should we care about the manipulation in bitcoin that took place in 2013? After all, the Bitcoin ecosystem is not nearly as important as the New York Stock Exchange. Nonetheless, recent trends indicate that bitcoin is becoming an important online currency and payment system. So it is important to understand how the Bitcoin ecosystem works or does not.

Additionally, trading in cryptocurrency assets has exploded recently. In the case of bitcoin, during the one year period ending in mid-May 2017, the market capitalization increased

[1] For an in-depth overview of how the Bitcoin ecosystem works, see Böhme et al. [4].

[2] The Bitcoin ecosystem includes the core network for propagating transactions, the blockchain, and many intermediaries such as currency exchanges, mining pools and payment processors that facilitate trade. We use "Bitcoin" with a capital "B" to refer to the ecosystem and "bitcoin " with a small "b" or BTC to refer to the coin.

massively, from around 7 Billion USD to 28 Billion USD [6]. That is an increase of approximately 300 percent in one year. The market cap of other cryptocurrencies surged by even more. In the one year period ending in mid-May 2017, the market value of cryptocurrencies excluding bitcoin surged from 1.7 Billion USD to more than 29 Billion USD [7]. That is an increase of more than 1,900 percent. Similar to the bitcoin market in 2013 (the period we examine), markets for these other crytocurrencies are very thin. Our analysis suggests that manipulation is quite feasible in such settings.

Trading in crytocurrencies are done over-the-counter (OTC). Such trades occur directly between two parties, that is, without going through a regulated stock exchange. OTC trading has exploded in recent years. In 2008 around 16 percent of U.S. stock trades were of the OTC type. By 2014, OTC trades accounted for forty percent of all stock trades in the US. Like cryptocurrency trading, OTC trades are not transparent and not regulated, and there is concern that such trading is more harmful than high-frequency trading via regulated exchanges [13].

As mainstream finance invests in cryptocurrency assets and as countries take steps toward legalizing bitcoin as a payment system (as Japan did in April 2017), it is important to understand how susceptible cryptocurrency markets are to manipulation. As this paper will show, the first time Bitcoin reached an exchange rate of more than $1,000, the rise was driven by fraud. It took more than three years for these exchange rates to be reached again, and we are left to wonder whether the current spike was driven by legitimate interest or by something more nefarious.

For all of these reasons, we believe that it is important to understand how the Bitcoin ecosystem works and how it could be abused. In this paper, we have taken an initial step in that direction.

1.2 Road Map

The paper proceeds as follows. Section 2 discusses background and related work. In section 3, we explain our methodology for identifying the STA and detail evidence for why we deem these transactions suspicious. Sections 4 and 5 examine the data in detail, present our findings and show that our results are robust. Section 6 has further discussion and brief conclusions.

2 Background and Related Work

Cryptocurrencies and associated markets represent a nascent but growing force within the financial sector. Bitcoin, which became the first popular decentralized cryptocurrency in 2009, is the most researched because it is the most successful of the digital currencies. Within the finance literature, there is growing interest in discovering what drives a "value-less" currency. Li and Wang investigate the bitcoin exchange rate in an effort to expand our understanding of the motivation behind the rise and fall of cryptocurrency values [12]. Additionally, Hayes constructs a model for determining the value of a bitcoin-like cryptocurrency by calculating its cost of production [10]. Ciaian et al. concluded that investor attractiveness has had a significant impact on bitcoin's price [18]. [3] While the potential for manipulation to influence valuations is sometimes acknowledged, none of these papers considered how unauthorized trades like the ones we study could affect valuations.

Unregulated cryptocurrency exchanges, such as Mt. Gox, are an essential part of the Bitcoin ecosystem. For most users, it is through currency exchanges that bitcoins are first acquired. As exhibited by the rise and fall of Mt. Gox, no cryptocurrency exchange is too big to fail. As reported by Moore and Christin, by early 2013 45% of Bitcoin exchanges had closed, and many of the remaining markets were subject to frequent outages and security breaches [15]. Vasek et al. performed an in-depth investigation of denial-of-service attacks against cryptocurrency exchanges and other Bitcoin services, documenting 58 such attacks [21]. Feder et. al [8] conducted the first econometric study of the impact of denial-of-service attacks on trading activity at Bitcoin exchanges, leveraging Vasek et al.'s data on attacks. They rely on the same leaked Mt. Gox trading data used in our paper, but use it to show how trading volume becomes less skewed (fewer large trades) the day after denial-of-service attacks targeted the Mt. Gox exchange. In this paper, we use the trading data to identify unauthorized trading and examine the effects of such trading on the Bitcoin ecosystem.

Due to their relatively lawless nature, cryptocurrencies are under constant threat of attack. Numerous researchers have conducted studies in order to document and combat threats such as Ponzi schemes [21], money laundering [16], mining botnets [11], and the theft of "brain" wallets [20]. Ron and Shamir attempt to identify suspicious trading activity by building a graph of Bitcoin transactions found in the public ledger [19].[4] None of these

[3]Gandal and Halaburda [9] examine competition among cryptocurrencies. They find that the data are consistent with strong network effects and winner-take-all dynamics.

[4]Meiklejohn et al. examine the blockchain to determine whether bitcoin transactions are are truly anonymous. They successfully link transactions back to popular Bitcoin service providers, such as currency exchanges [14].

papers can associate individual transactions with specific users at currency exchanges. Our data includes the user ID. Hence, we can associate trades with particular users.

For a more complete review of the literature, see Bonneau et al. for coverage of technical issues [5] and Böhme et al. for a discussion of Bitcoin's design, risks and open challenges [4].

3 Identifying Suspicious Trading Activity on Mt. Gox

3.1 Exchange Activity

In early 2014, in the midst of theft allegations, the Mt. Gox transaction history was leaked in the form of several large CSV files. The Mt. Gox data dump gave access to approximately 18 million matching buy and sell transactions which span April 2011 to November 2013. These data are much more finely grained than data we would be able to get from the blockchain for two reasons. First, a majority of the trading activity is recorded only by the exchange. Second, the exchange links transactions by the user account.

Data from the dump include fields such as transaction ID, amount, time, currency, and user country and state codes. Also included is the user ID, which is the internal number associated with Mt. Gox users. The user ID is crucial as it enables us to link transactions by the same actor.

We supplemented the Mt. Gox data with publicly available daily aggregate values from bitcoincharts.com. This data was used to verify trading volumes, to compare Mt. Gox exchange rates to other leading platforms, and to verify daily USD to BTC exchange rates.

3.2 Dataset Validation

With the exception of a few key steps, validating the Mt. Gox data closely followed previous work done by Feder et al. [8] in which duplicates were removed by inspecting combinations of key fields. The duplicate rows contained matching values for user ID, time stamp, transaction type (buy/sell), and transaction amount. We examined two methods to remove duplicate entries. Both methods treated tuples as unique (user ID, timestamp, transaction type, amount in BTC, amount in JPY, i.e., Japanese Yen) with the more aggressive of the two methods excluding amount in JPY from the tuple.[5] Both methods produced results that were more consistent with other publicly available trading data than the original dataset. Feder et al. [8] chose to proceed with the less aggressive of the two strategies, which resulted in a clean dataset of approximately 14 million records. We chose the more aggressive method,

[5] Mt Gox was based in Tokyo.

but our results are robust to both methods.

During the data exploration phase, we discovered additional duplicate records that did not fit the unique tuple model outlined above. In these instances they appeared to be copies of either one side (buy/sell) of the transaction or of the entire transaction with minor alterations to the data in the "User_ID," "Money," and "Money_JPY" columns. The common factor used to start the removal process was the new user ID. We could find one side of the transaction by matching on this user ID, and then use the Money and Money_JPY columns to find the matching opposite side of the transaction. In total 5,991 additional rows were removed using this method, all involving a single user ID. We later identified these duplicate entries as originating from the trader denoted "Markus." We performed additional sanity checks of the data utilizing publicly available historical Mt. Gox trading data from `bitcoincharts.org`. We are confident that the data are high-quality.

3.3 Suspicious Trading Activity

In early 2014, after the Mt. Gox data leak, several individuals trading on Mt. Gox found what they considered "suspicious activity" and wrote extensively about it [1, 3]. We conducted our own analysis of the data, confirming much of what was reported on the blogs. The rest of this section summarizes the evidence for why the trading activity should be deemed suspicious, along with a description of the behavior.[6] In Appendix B, we carefully go through the details that confirmed that the relevant transactions were suspicious. We present a summary of the key findings here.

3.3.1 Suspicious Trader 1: Markus

During initial data exploration we found a group of users with attributes that differed from the rest of the users in the dataset. In particular, for these users every transaction had "??" as an entry for the user country and user state fields. This appeared suspicious as these fields normally contain FIPS location codes, a NULL value, or "!!". One account containing the abnormal location values stood out when compared to the others because this account bought and sold bitcoins, where as the others only bought. We adhere to the naming convention in the blogs and refer to the first account as Markus.

Upon closer inspection, Markus's trades raised many red flags. He never paid transaction fees and reportedly paid seemingly random prices for bitcoins. Most curious of all, we identified many duplicate transactions in which the amount paid was changed from an

[6] Online commentary about these trades frequently refer to the traders as 'bots' (e.g., [1, 3]). We refrain from doing so in this paper because we have no evidence for whether or not the trades were issued automatically, as would be the case for bots.

implausibly random price to one that was consistent with other trades that day. In the end, we have concluded that Markus did not actually pay for the bitcoins he acquired; rather, his account was fraudulently credited with claimed bitcoins that almost certainly were not backed by real coins. Furthermore, because transactions were duplicated, no legitimate Mt. Gox customer received the fiat currency Markus supposedly paid to acquire the coins.

Markus began buying bitcoin on 2013-02-14 and was active until 2013-09-27. He did not pay for the bitcoins he acquired nor did he pay fees for the transactions. During the 225 days the account was active, Markus acquired a total of 335,898 bitcoins (worth around $ 76 Million) on 33 days.

3.3.2 Suspicious Trader 2: Willy

The remaining accounts found to contain "??" in the user state and user country fields were grouped separately from the Markus account because their trading activity looked different. Again, we adhere to the naming conventions found online [3, 1] and refer to this collection of accounts as "Willy". Unlike Markus, Willy did not use a single ID; instead, it was a collection of 49 separate accounts that each rapidly bought exactly 2.5 million USD in sequential order and never sold the acquired bitcoin. The first Willy account became active on 2013-09-27, a mere 7 hours and 25 minutes after Markus became permanently inactive, and we are able to track Willy activity until our data cutoff on 2013-11-30. Each account proceeded to spend exactly 2.5 million USD then it became inactive. Afterwards, the next account would become active and the process would repeat.

Why do we suspect foul play? Unlike Markus, there was no evidence of a cover up by introducing manipulated duplicate records. Indeed, it appears that the users who sold to Willy were in fact aware of the transaction. In addition to the circumstantial evidence of sequential use and proximity to Markus, the most solid evidence we have that foul play was involved can be traced to the internal user ID. Previous research into the account IDs used for this activity showed that they were abnormally high for the time period in which they operated [3]. Normal accounts for this time period had IDs that capped around 650000 where the users at the center of this research had IDs in the range of 658152-832432.

Furthermore, several reports can be found online of the Mt. Gox trading API going offline for various periods of time in which no trading activity was being processed with one exception; Willy trading activity continued unabashed [1]. On 2014-01-07 the trading API was offline for 90 minutes. During this time period the only activity being processed followed the exact buying pattern of Willy when he was active: 10-19 bitcoins purchased every 6-20 minutes.

In the 65 days between the first Willy transaction and our data cutoff, Willy trades were

made on 50 of those days. In total Willy bought around 268,132 bitcoin for just under $112 million. The number of bitcoins acquired by Willy was slightly less than the number of bitcoins that Markus acquired. However, the Markus activity occurred on 36 days over a 225 day period. Thus, the Willy activity was much more intense. Unlike, Markus, it appears that Willy was interacting with real users. While accounts of these users were "nominally" credited with Fiat currency, Willy did not actually pay for the bitcoins.

Hence, together, these unauthorized traders "acquired" around 600,000 bitcoins by November 2013. Perhaps unsurprisingly, this is very close to the number of bitcoins (650,000) that Mt. Gox claimed to have lost when it folded in early 2014.[7]

Theories on what motivated Willy's behavior Why did Willy purchase large quantities of Bitcoin? Was there a profit motive? If so, how did the ruse work? We cannot know for certain, but we will focus on two plausible explanations.

The first theory, initially espoused in a Reddit post shortly after Mt. Gox's collapse [2], is that hackers stole a huge number of BTC from Mt. Gox in June 2011 and that founder Mark Karpales took extraordinary steps to cover up the loss for several years.

Note that Bitcoin currency exchanges function in many ways like banks. Customers buy and sell bitcoins, but typically maintain balances of both fiat currencies and bitcoins on the exchange without retaining direct access to the currency. For bitcoin, this means that a customer account might reflect a balance of, say 100 BTC, but the customer does not retain access to the private keys that would enable her to spend the bitcoins directly. Instead, the user would have to request the keys to do so, which is analogous to a bank customer withdrawing cash from a local branch. Just as a bank maintains cash reserves that represent a fraction of total deposits, so too could an exchange represent to customers that they have more bitcoin in their accounts than is on hand.

If Mt. Gox was trying to hide the absence of a huge number of BTC from its coffers, it could succeed so long as customers remained confident in the exchange. By offering to buy large numbers of bitcoins, Willy could prop up the trading volume at Mt. Gox and "convert" consumer "bitcoin" balances to fiat money. The ruse could work so long as the users who sold bitcoin had enough confidence to leave the bulk of their fiat balance at Mt. Gox. Consequently, this theory holds that Willy was not trying to profit directly from these large purchases, but rather was trying to stave off collapse of the exchange.

This strategy would also be helpful even if some consumers requested to take our their fiat balance at Mt. Gox. If consumers wanted to take out bitcoins, Mt. Gox would immediately

[7] When Mt. Gox folded, it initially announced that around 850,000 bitcoins belonging to customers and the company were missing and likely stolen. Shortly thereafter, Mt. Gox found an additional 200,000 bitcoins. Hence, the total loss was 650,000 bitcoins.

have to supply them. On the other hand, if consumers wanted to redeem the fiat cash in their accounts, Mt. Gox could "delay" the withdrawal by saying that their bank was placing limits on how much fiat cash Mt. Gox could withdraw in a particular period. This indeed happened, and many consumers could not withdraw cash from their fiat accounts during the last couple of months before Mt. Gox shut down. By using this strategy, the trader turned the Mt. Gox' "bitcoin deficit" into a fiat currency deficit. This delayed the inevitable crash of Mt. Gox. Although this cannot work in the long-term, Bernie Madoff, a once respected stockbroker, kept a similar scheme running for many years.

As we will see in the subsequent sections, Willy's actions did produce a significant positive (albeit shortlived) effect on the price of bitcoin at Mt. Gox and other exchanges. This leads to a second, less conspiratorial, theory for the motivation behind Willy's behavior. Suppose that the user behind Willy had previously acquired bitcoins at a lower price. Many early adopters of bitcoin had acquired vast quantities at low prices. If one such user realized that due to a security weakness at Mt. Gox, there was a way to initiate costless bitcoin purchases (since the users who sold to Willy only received notional balances in their Gox account,) this user could then initiate these bulk purchases to drive up the exchange rate. Then the user could sell the bitcoins at a significant profit, either on Mt. Gox or on one of the other exchanges.

We do not know for sure which, if either, of these scenarios reflect what actually happened. But that is largely beside the point. Our goal is to demonstrate that these fraudulent trades did in fact significantly impact the price of bitcoin.

4 Impact of Suspicious Trading Activity: Preliminary Analysis

Figure 1 shows that the USD/BTC exchange rate increased dramatically during the period Willy was active. We are, of course, not the first to notice that. But that in itself does not mean that Willy's activity *caused* the price rise. In this section and the next, we provide compelling evidence that Willy's activity indeed *caused* the price rise.

In the unauthorized activity on Mt. Gox, Markus and Willy were almost always the buyers in the transactions. On the days they were active, they purchased large amounts of bitcoins.

As Table 1 shows, Markus was active as a buyer on 33 days, whereas Willy was active on 50 days. On the days Markus was active, he purchased on average $9,302$ BTC, which accounted for approximately 21 percent of Mt.Gox's daily volume of trades. On the days

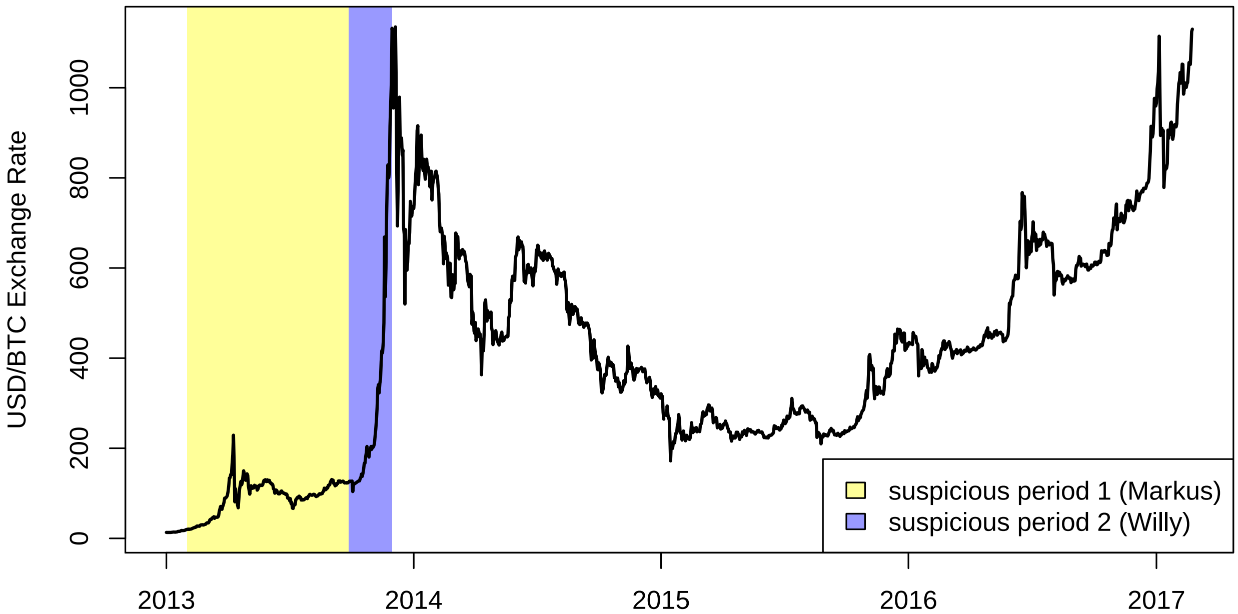

Figure 1: Bitcoin-USD exchange rate at Bitstamp exchange, with periods of suspicious activity shaded.

Table 1: Daily BTC purchased by Markus and Willy on days they were active.

	Mean	SD	Median	Min	Max	N
Markus:						
BTC purchased	9,302	7,310	5,874	696	24,785	33
% of Mt.Gox daily trade	0.21		0.17			
% of total trade	0.12		0.1			
Willy:						
BTC purchased	4,962	4,462	3,881	82	26,693	50
% of Mt.Gox daily trade	0.18		0.15			
% of total trade	0.06		0.05			

Willy was active, he purchased on average $4,962$ BTC, which accounted for 18 percent of Mt.Gox's daily volume of trades. In both cases these are substantial amounts. The percent of the BOTs' purchasing was also a non-trivial amount of the total trade in bitcoins, as the Table shows. Marcus accounted for 12 percent and Willy accounted for 6 percent of the total trade on the four main exchanges trading bitcoin and USD on the days they were active. In addition to Mt. Gox, the other main exchanges trading USD/BTC during this time period were Bitstamp, Bitfinex and BTC-e. [8]

[8] See Appendix D for the market share of the exchanges.

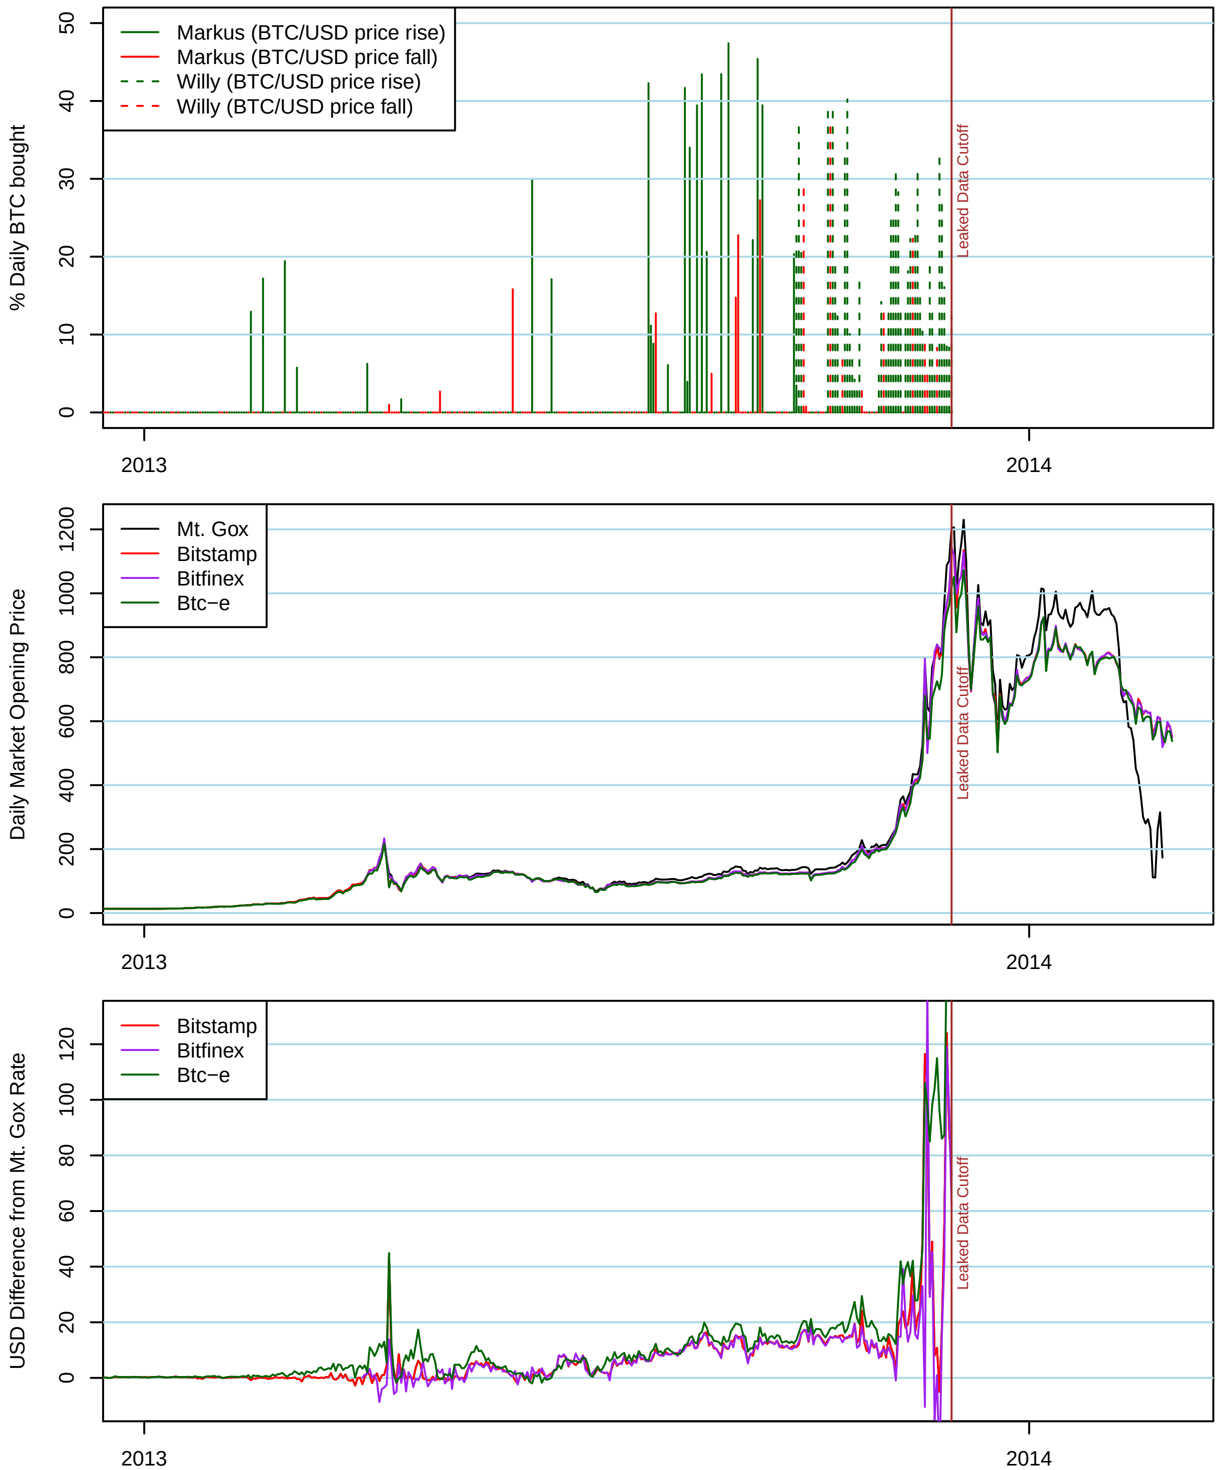

Figure 2: Top: Percentage of total daily trade volume at Mt. Gox when Willy and Markus are active; shaded green if the BTC/USD exchange rate closed higher and red otherwise. Center: BTC/USD exchange rate over time at Mt. Gox and other leading exchanges. Bottom: Difference in exchange rate between Mt. Gox and other leading exchanges.

4.1 Suspicious Purchases and Price Changes

We would expect an association between the suspicious purchases and a rise in prices on Mt. Gox (and other exchanges as well.) This is because an upwards shift in demand should lead to a rise in price. We will examine this issue in this section. Although the activity took place exclusively on Mt. Gox, we are also interested in examining how it affected the other exchanges in the Bitcoin ecosystem.

The top graph in Figure 2 shows the fraction of daily BTC traded on the Mt. Gox exchange platform that were carried out by Markus and Willy, respectively. It is important to note that Markus and Willy's activity overlapped by one day only. On the days that there was suspicious trading activity on Mt. Gox, the descriptive evidence suggests that prices also tended to rise. The lines in the figure are colored green if the exchange rate rose and red if the exchange rate fell.

We then examined whether the price changes differed on the days in which the fraudulent activity occurred. We did this for the 9.5 months Markus and Willy were active (and for which we have data) and observed how often the exchange rate rose on Mt. Gox, as indicated in Table 2. We can see that on days without suspicious activity, 55% of the time the exchange rate did in fact rise. But on the 82 days that there was suspicious purchasing activity, 79% of the time the exchange rate rose. According to a chi-squared test of proportions, it is unlikely that this difference was due to randomness ($p < 0.05$). This is preliminary evidence that this activity did contribute to price rises on Mt. Gox.

Table 2: Unauthorized activity and price changes on Mt. Gox

		Days with no STA		Days with STA	
		days	%	Days	%
Markus	Daily rate decrease	84	44	7	21
	Daily rate increase	109	56	26	79
Willy	Daily rate decrease	9	60	10	20
	Daily rate increase	6	40	40	80
Total	Daily rate decrease	93	45	17	21
	Daily rate increase	115	55	65	79

Not surprisingly, Markus and Willy's activity affected the prices on the other platforms as well. As shown in Table 3, on days without unauthorized activity, the exchange rate on Bitstamp rose 55% of the time. However, on the 82 days that Markus or Willy acquired bitcoins, the exchange rate rose more than 80 percent of the time. Hence there is also strong

evidence that the effects of suspicious trades on Mt. Gox spilled over to other exchanges. This makes sense because all of these platforms traded the same USD-BTC currency pair.

Table 3: Suspicious trading activity and price changes on Bitstamp

		Days with no STA		Days with STA	
		days	%	Days	%
Markus	Daily rate decrease	88	45	6	18
	Daily rate increase	105	55	27	82
Willy	Daily rate decrease	6	40	9	18
	Daily rate increase	9	60	41	82
Total	Daily rate decrease	94	45	15	18
	Daily rate increase	114	55	67	82

We then divided the data into five equal three-month periods, starting from 2012-12-01 (2.5 months before Markus was active) and ending the date Mt. Gox ceased operations (2014-25-02,) which was three months after the leaked Mt.Gox dataset ends. For the fifth period we do not have any data regarding suspicious trading activity activity, but we do have data regarding prices on Mt.Gox and the other platforms.

Table 4: Suspicious trading activity: % of days active during each period

	Period 1 2012-12-01 – 2013-02-28	Period 2 2013-03-01 – 2013-05-31	Period 3 2013-06-01 – 2013-08-31	Period 4 2013-09-01 – 2013-11-30	Period 5 2013-12-01 2013-02-25
Markus	3%	5%	19%	9%	no data
Willy	0	0	0	55%	no data
N	90	92	92	91	87

Table 4 shows the percent of days in each period, in which there was suspicious trading activity. Markus was active over 8 months, which span over 4 periods. However, he was primarily active in period 3. Willy on the other hand was active for less than three months and all of the activity occurred in period 4. We have no data on any unauthorized activity from the end of period 4. Mt. Gox shut down shortly thereafter.

In Table 5 we see how the daily exchange rate (closing − opening) changed, on average, on 4 exchange platforms. Since fraudulent activity essentially only occurred in the third and fourth periods, we focus on these two periods. But periods one and two can be viewed as benchmarks.

Table 5: Average daily rate changes in USD-BTC exchange rate by period in $

	Period 1	Period 2	Period 3			Period 4		
			All	Markus active	Markus not active	All	Willy active	Willy not active
Rate change Mt.Gox	0.21	1.00	0.16	3.15	-0.51	11.61	21.85	-0.88
Rate change Bitstamp	0.23	1.02	0.02	2.35	-0.51	10.99	20.37	-0.45
Rate change Bitfinex	.	0.92	0.04	2.14	-0.44	10.75	19.54	0.03
Rate change Btce	0.22	1.05	-0.1	1.81	-0.53	10.30	19.22	-0.58
N	90	92	92	17	75	91	50	41

In period 3, when Markus' activity peaked, we don't see much change overall in prices. However, if we look at the days Markus is active, the average daily price increase is higher. This is true, both on Mt. Gox and on all the other platforms too.

In period 4, the sole period in which Willy was active, we see a big jump in the average daily exchange rate change. Separating the days on which Willy was active from those he was not, reveals a dramatic difference: In the case of Mt. Gox, the average USD/BTC rate increased by $21.85 on the 50 days Willy was active; it actually fell (by $0.88 on average) on days when Willy was not active.

The same dramatic difference holds for the other exchanges as well. These results are striking and suggest that Willy's activity could have *caused* huge jumps in the exchange rate on all of the exchanges. We will run regressions to control for other variables in Section 5, but these summary statistics make it very clear that the suspicious purchasing activity likely caused the huge price increases.[9]

4.2 Suspicious Activity and Price Differences Between Exchanges

Another interesting price indicator, is the difference in opening prices on Mt. Gox and other platforms. In the middle graph in Figure 2 we plot the differences in daily opening prices

[9]In period 4, Willy was active on 50 out of 91 days. But since Willy does not begin trading until after Markus ceases activity, all of Willy's activity takes place in a 65 day window in Period 4 from September 27, 2013 until October 30, 2013 (the end of period 4.) He is active for 50 of the sixty five days. In Appendix C, we show Table 5 for this sub-period of period 4. The "patterns" shown in Table 5 are qualitatively unchanged.

(in percentage terms) between Mt. Gox and the other three platforms.

The graph shows that there were differences in opening prices between Mt. Gox and the other exchanges, and these differences took take place when the suspicious activity occurred.

Table 6: Average percentage difference in opening prices

	Period 1	Period 2	Period 3	Period 4	Period 5
Mt.Gox - Bitstamp	1.0	1.6	6.3	8.6	0.3
Mt.Gox - Bitfinex	.	1.0	6.5	8.1	0.3
Mt.Gox - Btce	2.2	5.6	7.6	12.2	1.5
N	90	92	92	91	87

To explore this further, we divide the data into 5 periods. In Table 6 we have the average percentage difference in daily "opening" prices between Mt. Gox and the other platforms. This measures by how much percent was Mt. Gox's opening price higher (or lower) than the other platforms' prices.[10]

We see an interesting trend in the price differences. In periods one and two (before the significant fraudulent activity), there is relatively little difference between Mt. Gox and the other exchanges. Similarly in period 5, there is again very little difference in prices between Mt. Gox and the other exchanges.

The percentage differences grow in periods periods 3 and 4, the periods in which STA is most prevalent. This trend reverses in period 5, which is the period in which Mt. Gox shut down. Although we do not have data on suspicious trading activity during this period, the small opening price differences between Mt. Gox and the other exchanges suggests that these players were not active then or were less active. This suggests that during the time there was significant fraudulent activity, the activity had a role in creating a gap between Mt.Gox and the other platforms.

[10]The same is true for the closing prices. Closing prices on day t equal opening prices of day $t + 1$ since there is 24 hour trading. The opening/closing price is at 24:00 GMT.

5 Regression Analysis

The analysis in the previous section provide strong evidence that the suspicious activity on Mt. Gox affected prices on all exchanges. To further examine this, we use regression analysis, although given the clear results from the summary statistics in section 4, the regression analysis is probably not necessary. In any case, we run regressions with the dependent variable being the price change on Mt. Gox. We then run the same regressions for the other three exchanges as well.

5.1 Fluctuations of Prices on Mt. Gox

We run the following regression:

$$RateChange_t = \beta_0 + \beta_1 Markus_t + \beta_2 Willy_t + \beta_3 DDoS_t + \beta_4 DayAfterDDoS_t + \beta_5 Other_t + \epsilon_t \tag{1}$$

Our dependent variable, *RateChange*, is the daily difference in the exchange rate of BTC, i.e. the daily difference between the closing and opening prices on Mt.Gox.[11]

We now describe our independent variables. *Markus* is a dummy variable that takes on the value one on the days Markus is active as a buyer. Similarly, we define the dummy variable *Willy*. *DDoS* is a dummy variable that takes on the value one on days a DDoS attack on Mt. Gox occurred. *Day after DDoS* is a dummy variable that takes on the value one on the day after a DDoS attack on Mt. Gox occurred. The variable *Other* (or *OtherAttacks*) is a dummy variable that takes on the value one on days that non DDoS attacks occurred.[12] ϵ_t is a white noise error term.[13] The subscript "t" refers to time. In the analysis, we used data from the first four periods, since in period 5, we have no data on whether the fraudulent traders were active. We have a total of 365 observations.

Equation (1) is a reduced-form regression. That is, we are not estimating demand or supply, but rather the effect of changes in exogenous right-hand-side variables on the endogenous variable (daily price change.) But in our case, the coefficients from this reduced form regression are exactly what we want to measure. Summary statistics appear in Appendix A.

[11] Recall that closing prices on day t equal opening prices of day $t+1$ since there is 24 hour trading. The opening/closing price is at 24:00 GMT.

[12] Perhaps because it was the leading exchange during the period of our data, most of the DDoS attacks were on Mt. Gox.

[13] We check for autocorrelation of errors by calculating the Durbin Watson (DW) statistic for each regression. The value of DW is not statistically different from two in any of the four cases; this strongly suggests that there is no autocorrelation and a white noise error term is appropriate.

5.1.1 Results

Table 7: Examining Price Changes Within Mt. Gox and the other platforms

Independent Variables / Dependent Variable	Mt.Gox Rate Change	Bitstamp Rate Change	Bitfinex Rate Change	BTC-E Rate Change
Markus	2.79	3.24	2.06	2.37
	(0.72)	(0.96)	(0.31)	(0.71)
Willy	21.65***	20.21***	19.23***	19.04***
	(6.66)	(7.18)	(3.63)	(6.81)
DDoS	-2.38	-1.67	-1.87	-2.01
	(-0.55)	(-0.44)	(-0.26)	(-0.54)
Day After DDoS	-3.50	-3.25	-2.9	-2.68
	(-0.80)	(-0.86)	(-0.41)	(-0.72)
Other Attacks	7.16	5.70	7.35	5.61
	(0.82)	(0.75)	(0.44)	(0.75)
N	365	365	244	365
adj. R^2	0.104	0.120	0.037	0.108

t statistics in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

The results in Table 7 show that the coefficient representing Willy's activity is positive and significant: hence there is a very strong positive association between activity by Willy and prices on Mt. Gox. This regression confirms the striking results of Section 4. The estimated coefficient on the "dummy" variable for Willy is $21.65, while the "estimate" in section 4 was $21.85. This again suggests that the USD/BTC exchange rate rose on Mt. Gox by more than 20 dollars a day on average on the days that Willy was active.

The regressions for the other exchanges in the same table shows that price on that exchange also rose by 19-20 dollars a day on average on the days that Willy was active. Again the estimated coefficients are consistent with the "estimates" from Table 5 in section 4.

Note that for these regressions, the estimated coefficient on the dummy variable representing Willy's activity is the only coefficient which is significant. Notably, denial-of-service

attacks and other shocks do not appear to influence the exchange rate.

The estimated coefficient associated with Markus's activity is positive, but not significant, suggesting that Willy's intense activity had more of an effect on prices than did Markus' more diffused activity.

6 Conclusion

In this paper, we used trade data delineated by user to determine whether there was (as claimed in the popular press) suspicious trading activity on the Mt. Gox exchange. We find overwhelming evidence of suspicious/fraudulent activity on Mt. Gox. We then showed how this activity affected the Bitcoin ecosystem.

We have shown that manipulations can have important real effects. The suspicious trading activity of a single actor caused the massive spike in the USD-BTC exchange rate to rise from around \$150 to over \$1 000 in late 2013. The fall was even more dramatic and rapid, and it has taken more than three years for Bitcoin to match the rise prompted by fraudulent transactions.

Because such delineated data are virtually never available to researchers, we cannot say whether such activity continues to plague the Bitcoin ecosystem. Given the recent meteoric rise in bitcoin to levels beyond the peak 2013 (and the huge increase in the prices of other cryptocurrencies), it is important for the exchanges to ensure that there is not fraudulent trading. Since the Bitcoin ecosystem is currently unregulated, "self-policing" by the key players and organizations is essential. Additionally, regulators may want to begin taking an active oversight role as the Bitcoin ecosystem becomes more integrated into international finance and payment systems.

References

[1] Free willy! – identifying the gox buy bot., January 2014. https://www.reddit.com/r/Bitcoin/comments/20k4zc/free_willy_identifying_the_gox_buy_bot/.

[2] Peter Rs theory on the collapse of Mt. Gox, March 2014. https://www.reddit.com/r/Bitcoin/comments/1zdnop/peter_rs_theory_on_the_collapse_of_mt_gox/.

[3] The Willy Report, May 2014. https://willyreport.wordpress.com/.

[4] Rainer Böhme, Nicolas Christin, Benjamin Edelman, and Tyler Moore. Bitcoin: Economics, technology, and governance. *Journal of Economic Perspectives*, 29(2):213–38, 2015.

[5] Joseph Bonneau, Andrew Miller, Jeremy Clark, Arvind Narayanan, Joshua A. Kroll, and Edward W. Felten. Research Perspectives and Challenges for Bitcoin and Cryptocurrencies. In *IEEE Symposium on Security and Privacy*, May 2015.

[6] CoinMarketCap. Cryptocurrency market capitalizations. https://coinmarketcap.com/currencies/bitcoin/. Last accessed May 16, 2017.

[7] CoinMarketCap. Total market capitalization (excluding bitcoin). https://coinmarketcap.com/currencies/bitcoin/. Last accessed May 16, 2017.

[8] Amir Feder, Neil Gandal, JT Hamrick, and Tyler Moore. The impact of DDoS and other security shocks on Bitcoin currency exchanges: Evidence from Mt. Gox. In *15th Workshop on the Economics of Information Security (WEIS)*, 2016.

[9] Neil Gandal and Hanna Halaburda. Can we predict the winner in a market with network effects? In *Games*, July 2016.

[10] Adam S. Hayes. Cryptocurrency value formation: An empirical study leading to a cost of production model for valuing bitcoin. *Telematics and Informatics*, May 2016.

[11] Danny Yuxing Huang, Hitesh Dharmdasani, Sarah Meiklejohn, Vacha Dave, Chris Grier, Damon McCoy, Stefan Savage, Nicholas Weaver, Alex C Snoeren, and Kirill Levchenko. Botcoin: Monetizing stolen cycles. In *Proceedings of the Network and Distributed System Security Symposium*, 2014.

[12] Xin Li and Chong Alex Wang. The technology and economic determinants of cryptocurrency exchange rates: The case of bitcoin. *Decision Support Systems*, December 2016.

[13] John McCrank. Dark markets may be more harmful than high-frequency trading. *Reuters Business News*, apr 2014. http://www.reuters.com/article/us-dark-markets-analysis-idUSBREA3508V20140406.

[14] Sarah Meiklejohn, Marjori Pomarole, Grant Jordan, Kirill Levchenko, Damon McCoy, Geoffrey M Voelker, and Stefan Savage. A fistful of Bitcoins: Characterizing payments among men with no names. In *Proceedings of the Internet Measurement Conference*, pages 127–140. ACM, 2013.

[15] Tyler Moore and Nicolas Christin. Beware the middleman: Empirical analysis of Bitcoin-exchange risk. In *Financial Cryptography and Data Security*, volume 7859 of *Lecture Notes in Computer Science*, pages 25–33. Springer, April 2013.

[16] Malte Möser, Rainer Böhme, and Dominic Breuker. An inquiry into money laundering tools in the Bitcoin ecosystem. In *Proceedings of the Seventh APWG eCrime Researcher's Summit*, pages 1–14. IEEE, 2013.

[17] Satoshi Nakamoto. Bitcoin: A peer-to-peer electronic cash system. https://bitcoin.org/bitcoin.pdf, 2008.

[18] Miroslava Rajcaniova Pavel Ciaian and dArtis Kancs. The economics of bitcoin price formation. *Applied Economics*, 48:1799–1815, May 2016.

[19] Dorit Ron and Adi Shamir. Quantitative analysis of the full Bitcoin transaction graph. In *Financial Cryptography and Data Security*, volume 7859 of *Lecture Notes in Computer Science*, pages 6–24. Springer, 2013.

[20] Marie Vasek, Joseph Bonneau, Ryan Castellucci, Cameron Keith, and Tyler Moore. The Bitcoin brain drain: a short paper on the use and abuse of bitcoin brain wallets. In *Financial Cryptography and Data Security*, Lecture Notes in Computer Science. Springer, 2016.

[21] Marie Vasek and Tyler Moore. There's no free lunch, even using Bitcoin: Tracking the popularity and profits of virtual currency scams. In Rainer Böhme and Tatsuaki Okamoto, editors, *Financial Cryptography and Data Security*, volume 8975 of *Lecture Notes in Computer Science*, pages 44–61. Springer, January 2015.

Appendix A: Descriptive Statistics

Table 8: Summary statistics of independent and dependent variables

	Mean	SD	Min	Max
Markus	0.09	0.29	0	1
Willy	0.14	0.34	0	1
DDOS	0.08	0.27	0	1
Day after DDOS	0.08	0.27	0	1
Other Attacks	0.02	0.13	0	1
Mt.Gox rate change	3.24	22.39	-139.78	257.5
Bitstamp rate change	3.06	19.53	-132.99	190.91
Bitfinex rate change[14]	4.25	33.30	-295.97	294
Btce rate change	2.86	19.28	-134.30	198.14
N	365			

Table 9: Correlation between daily rate changes and the independent variables

	Mt.Gox Rate Change	Bitstamp Rate Change	Bitfinex Rate Change	Btce Rate Change
Markus	0.001	0.01	-0.02	0.00009
Willy	0.33	0.35	0.23	0.34
DDoS	-0.06	-0.06	-0.05	-0.06
Day After DDoS	-0.07	-0.07	-0.05	-0.06
Other Attacks	0.02	0.02	0.013	0.02
N	365	365	244	365

[14] N=244 for this variable.

Table 10: Correlation between independent variables

	Markus	Willy	DDoS	Day After DDoS	Other Attacks
Markus	1				
Willy	-0.1	1			
DDoS	0.05	-0.06	1		
Day After DDoS	0.05	-0.06	0.33	1	
Other Attacks	0.03	-0.05	-0.04	0.04	1
N	365				

Appendix B: Details of Markus and Willy Activity

Markus seemingly paid random rates for the bitcoins he acquired. For example, in two transactions that took place the same hour on 2013-04-03, he paid 0.000374 USD per bitcoin on one transaction and 925 489.67 USD per bitcoin on another.

Table 11 shows the wide range of rates that Markus paid. The table reports the number of purchases that Markus made for different ranges of rates. During the time when Markus traded, published exchange rates ranged from $20 to $229. Hence, any transactions with rates outside this range raise suspicion. In fact, only a quarter of Markus's trades fell within this range. 13% of the time, Markus paid less than one dollar, while in 821 transactions (3% of the total), he supposedly paid a rate of exceeding $100,000 per bitcoin!

Table 11: Distribution of USD/BTC rates paid by Markus

	$\leq \$0.10$	$> \$0.10$, $\leq \$1$	$> \$1$, $\leq \$20$	$> \$20$, $\leq \$229$	$> \$229$, $\leq \$1K$	$> \$1K$, $\leq \$10K$	$> \$10K$, $\leq \$100K$	$> \$100K$
#	1 050	2 586	6 320	7 009	3 658	4 604	2 311	821
%	3.7%	9.2%	22.3%	24.7%	12.9%	16.2%	8.1%	2.9%

Upon closer inspection, the random exchange rates appear to come from transactions posted before Markus' transactions. Table 12 illustrates the pattern. Transaction 1362466144485228 was posted with user 238168 buying ≈0.398 bitcoin for 15.13 USD. Every Markus transaction that followed (indicated in bold) "borrowed" the Money, and Money_JPY values from the previous transaction. We confirmed this pattern of behavior throughout – whenever Markus bought, the amount paid came from a previous unrelated transaction, while the number of

bitcoins acquired appears randomly.

Table 12: Fraudulent transactions initiated by Markus (user ID in bold)

Trade_Id	Date	User_Id	Type	Bitcoins	Money	Money_JPY
1362466099116388	2013-03-05 6:48	238168	buy	0.58932091	22.39419	2094.796
1362466099116388	2013-03-05 6:48	109955	sell	0.58932091	22.39419	2094.796
1362466144485228	2013-03-05 06:49	238168	buy	0.3982007	15.13163	1415.442
1362466144485228	2013-03-05 06:49	132909	sell	0.3982007	15.13163	1415.442
1362466154623847	2013-03-05 06:49	**698630**	buy	1.70382	15.13163	1415.442
1362466154623847	2013-03-05 06:49	96376	sell	1.70382	15.13163	1415.442
1362466154714939	2013-03-05 06:49	**698630**	buy	1	15.13163	1415.442
1362466154714939	2013-03-05 06:49	201601	sell	1	15.13163	1415.442

Occasionally Markus would also sell bitcoin, and the BTC-fiat currency exchange rate for these transactions appears to be correct. For example, on 2013-06-02 Markus sold 31.5 bitcoins for 3 757.95 USD, or 119.3 USD per bitcoin, which is similar to the average rate paid by users that day. In total, Markus sold 35867.18 bitcoin worth approximately 4 018 681.65 USD in 2927 transactions on 6 different days.

As stated in section 3.2, we paid closer attention to what records to remove while deduplicating the data. This was due to the fact that several transactions contained duplicate buy and sell rows; see Table 13 for an example of these transactions. We can see that apparently user 201601 sold one bitcoin twice at the same exact time, first to user 698630 for 15.13 USD and second to user 634 for 38.11 USD.

Table 13: Duplicate Markus Transactions

Trade_Id	Date	User_Id	Type	Bitcoins	Money	Money_JPY
1362466154714939	2013-03-05 06:49	201601	sell	1	15.13163	1415.442
1362466154714939	2013-03-05 06:49	**698630**	buy	**1**	**15.13163**	1415.442
1362466154714939	2013-03-05 06:49	201601	sell	1	38.11000	3564.883
1362466154714939	2013-03-05 06:49	**634**	buy	**1**	**38.11000**	3564.883

Upon closer inspection, we concluded that the rows containing 15.13163 in the money columns are the original rows for this transaction. In every instance where duplicates were discovered they involved user 698630 and user 634; 634 appeared to “correct” the 698630. There are multiple oddities involving user 634. First, the numeric user ID is extremely low, which strongly suggests that it could be an internal Mt. Gox account. Second, prior to issuing the corrected transactions, user 634 bought and sold a total of 824,297.7 bitcoin between 2011-04-07 and 2012-08-01. This account was inactive for 197 days before we see it used again in the duplicate transactions involving Markus.

Table 14 summarizes the discrepancies between Markus's identities. 2 966 buy transactions made by 698630 were later duplicated as originating from user 634 at market prices. In total, as user 698630, Markus reportedly paid 1 080 617 USD for 67 452 bitcoin. When acting as user 634 instead, Markus "paid" 2 000 729 USD for the same transactions. This only includes the corrected transactions involving user 634; we ignore any trading activity that occurred before Markus was active. It is worth noting that only the amounts paid for bitcoins were altered, never the bitcoin amount. Additionally, for the 196 transactions where user 698630 sold bitcoin and we found a duplicate row with user 634, no monetary amounts were altered. Only the user ID had changed.

Finally, it is worth noting that the majority of transactions by user 698630 were never changed, despite the presence of often wild exchange rates. User 698630 only operated between February and September 2013, and during that time he purchased 268 446.09 BTC, reportedly at prices totaling $76.4 million. We note that this total USD amount should be viewed with caution, given that it is based on seemingly random exchange rates.

Table 14: Summary of Markus transactions

	User ID	# Transactions	Total BTC	Total USD
Manipulated Buy	698630	2966	67 451.61	$1.1M
Manipulated Buy	634	2966	67 451.61	$2.0M
Unchanged Buy	698630	25407	268 446.09	$76.4M
Manipulated Sell	698630	196	5 049.86	$0.2M
Manipulated Sell	634	196	5 049.86	$0.2M
Unchanged Sell	698630	2 927	35 867.18	$4.0M

Appendix C: Further Examination of Period 4

Table 15: Price changes in the period Willy is active

	All days	Willy active	Willy not active
Mt. Gox	16.34	21.85	-2.006
Bitstamp	15.39	20.37	-1.237
Bitfinex	15.14	19.54	0.497
Btce	14.46	19.22	-1.411
N	65	50	15

Numbers in the Table are means
From our data, Willy was active from 27.9.2013
and until the end of the data, 30.11.2013

Table 16: Willy: Volume Activity - condensed 65 day period

	mean	sd	median	N
Volume bought by Willy	4,962	4,462	3,881	50
Total BTC volume (Willy active)	30,854	23,145	25,939	50
Total BTC volume (Willy inactive)	24,303	29,949	12,582	15

Table 17: Willy: Volume activity in all of period 4

	mean	sd	median	N
Volume bought by Willy	4,962	4,462	3,881	50
Total BTC volume (Willy active)	30,854	23,145	25,939	50
Total BTC volume (Willy inactive)	17,472	19,808	10,444	41

Appendix D: Bitcoin Marketshare

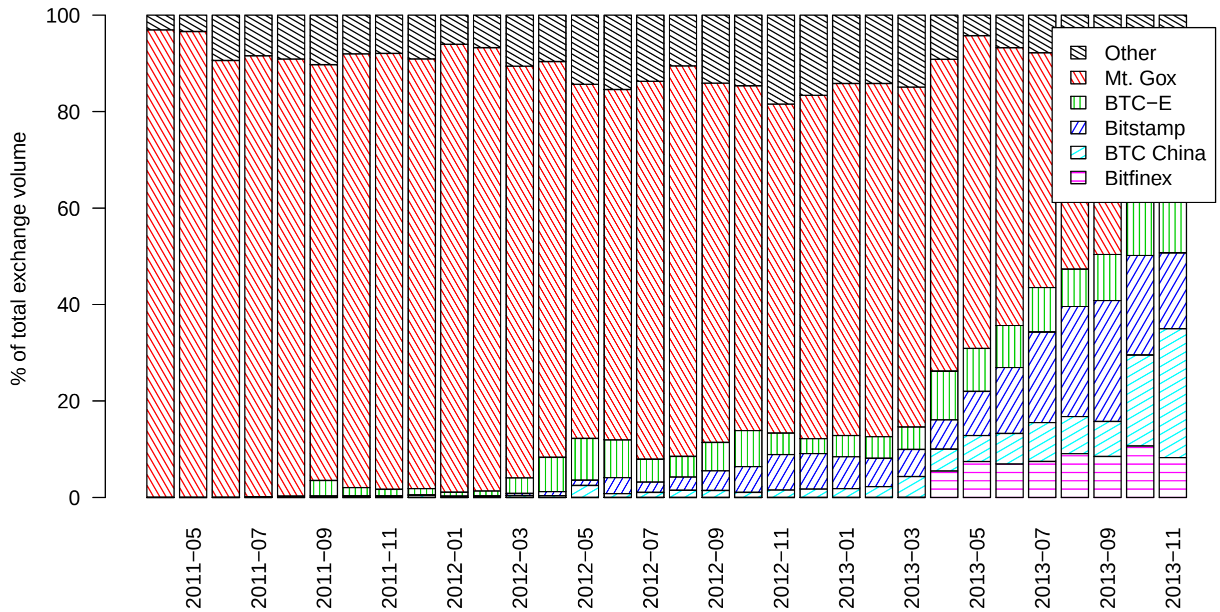

Figure 3: Distribution of market share among Bitcoin currency exchanges by reported trade volume, April 2011 to November 2013. (Source: `bitcoincharts.com`)